



AT6000 FPGAs

Application Note

3 x 3 Convolver with Run-time Reconfigurable Vector Multiplier in Atmel AT6000 FPGAs

Introduction

Convolution is one of the basic and most common operations in both analog and digital domain signal processing. Often times, it is desirable to modulate a given signal in order to enhance or extract important information contained in it. This signal modulation is generally known as filtering. In two-dimensional digital signal processing (2-D DSP), a 3 x 3 convolver is commonly used in applications such as object detection and feature extraction of 2-D images. However, due to the compute-intensive nature (multiply-add operations) of 2-D convolution, the size of the required digital logic circuitry increases accordingly. In this application note, we present an efficient single-chip FPGA implementation of a bit-parallel 3 x 3 symmetric convolver that features run-time software-configurable convolver coefficients (taps). Our reference design takes full advantage of time-division multiplexing (TDM), pipelining, CacheLogic[®], and vector multiplication. We will also discuss an alternative to using CacheLogic, which features run-time hardware self-configurable LUT using static random access memory (SRAM).

Convolution Basics

Unlike the “flip-and-drag” procedure performed in the analog-domain convolution, the 2-D digital-domain convolution is much easier to perform from a graphical standpoint. Given two 2-D infinite-impulse arrays $a(m,n)$ and $h(m,n)$, where m and n are integer indices, the convolution sum $p(m,n)$ is defined as:

$$p(m, n) = \sum_{i=-\infty}^{\infty} \sum_{j=-\infty}^{\infty} a(i, j) \cdot h(m-i, n-j) \quad [1]$$

Graphically, when we convolve $a(m,n)$ with $h(m,n)$, we multiply each tap in $a(m,n)$ with the corresponding tap in $h(m,n)$. The convolution output is then the sum of all multiplies. A 3 x 3 finite-impulse symmetric convolver example is shown below. $H(m,n)$ denotes the constant convolver mask. The underlined tap denotes the origin of an array.

Given:

$$a(m, n) = \begin{bmatrix} 1 & 2 & 4 \\ 0 & \underline{2} & 1 \\ 1 & 0 & 3 \end{bmatrix} \quad h(m, n) = \begin{bmatrix} 1 & 0 & 1 \\ 0 & \underline{2} & 0 \\ 1 & 0 & 1 \end{bmatrix}$$

Then,

$$p(0, 0) = \underline{13}$$

In the above example, our output origin has a value of $13 = (1 \times 1 + 2 \times 0 + 4 \times 1 + 0 \times 0 + 2 \times 2 + 1 \times 0 + 1 \times 1 + 0 \times 0 + 3 \times 1)$. Other entries of the output array will be computed as $a(m,n)$ slides across $h(m,n)$ [i.e., when the origin shifts in $a(m,n)$.]

An important property of a symmetrical 3 x 3 mask is that it allows for “pre-addition” of certain input values before any multiplication takes place. As in the above example, notice that we have only three constants in $h(m,n)$. Instead of performing nine multiplications, we can reduce the number to three by pre-adding the input values that get multiplied to each of the three mask constants. We get the same result of $13 = [(1 + 4 + 3 + 1) \times 1 + (2 + 1 + 0 + 0) \times 0 + 2 \times 2]$.



Although the number of multiplication is greatly reduced, a single-chip FPGA ($\cong 10K$ usable gates) implementation is impossible with fully pipelined multipliers due to their fairly large size. In view of this problem, vector multipliers have been proposed to provide an alternative to regular multipliers.

Vector Multiplier - Theory of Operation

Suppose we need to perform the following:

$$p = a(1) \cdot h(1) + a(2) \cdot h(2) + a(3) \cdot h(3) \quad [2]$$

where the $a(i)$ (input mask taps) and the $h(i)$ (constant convolver taps) are 2-bit positive integers. Let:

$$\begin{aligned} a(1) &= 01 & h(1) &= 11 \\ a(2) &= 11 & h(2) &= 00 \\ a(3) &= 10 & h(3) &= 10 \end{aligned}$$

By convention, the multiply-add operation takes place in a "top-down" fashion. First, we compute all stages of partial products of multiplying two numbers. Then we sum all partial products with appropriate arithmetic shifts to obtain each product. We then repeat the same procedure for all other multiplies. Finally we sum all products to get the output (see Figure 1 below).

Figure 1. Conventional Multiply-add Operation

Constants $h(n)$	11	00	10
Inputs $a(n)$	<u>x 01</u>	<u>x 11</u>	<u>x 10</u>
	11	00	00
	<u>00</u>	<u>00</u>	<u>10</u>
	011	000	100

$$P = 011 + 000 + 100 = 111$$

Table 1. Look-up Table

LSB Combination	Operation	Partial Product P
000	$0 + 0 + 0$	0
001	$0 + 0 + h(3)$	10
010	$0 + h(2) + 0$	0
011	$0 + h(2) + h(3)$	10
100	$h(1) + 0 + 0$	11
101	$h(1) + 0 + h(3)$	101
110	$h(1) + h(2) + 0$	11
111	$h(1) + h(2) + h(3)$	101

Another approach, which leads to the theory of vector multiplication using LUT (Look-up Tables), is to consider summing partial products of the same bit-level (see Figure 2). The output is then the sum of all partial product stages $P1, P2, \dots, Pn$ with the appropriate arithmetic shifts.

Figure 2. Alternative Multiply-add Operation

constants $h(n)$	11	00	10
inputs $a(n)$	<u>x 01</u>	<u>x 11</u>	<u>x 10</u>
partial prod. P1 =	11 +	00 +	00 = 11
partial prod. P2 =	<u>00</u> +	<u>00</u> +	<u>10</u> = 10

$$P = P1 + P2 = 11 + 100 = 111$$

Notice that when forming the partial products, we sum copies of the constant multiplicand $h(n)$. For example, when forming partial product P1, we bring down a copy of $h(n)$ when the least-significant bit (LSB) of the corresponding $a(n)$ is 1. If the LSB is 0, we simply bring down a zero. As a result, for the LSB combination of 110 (shown in Figure 2 as bold numbers), we perform:

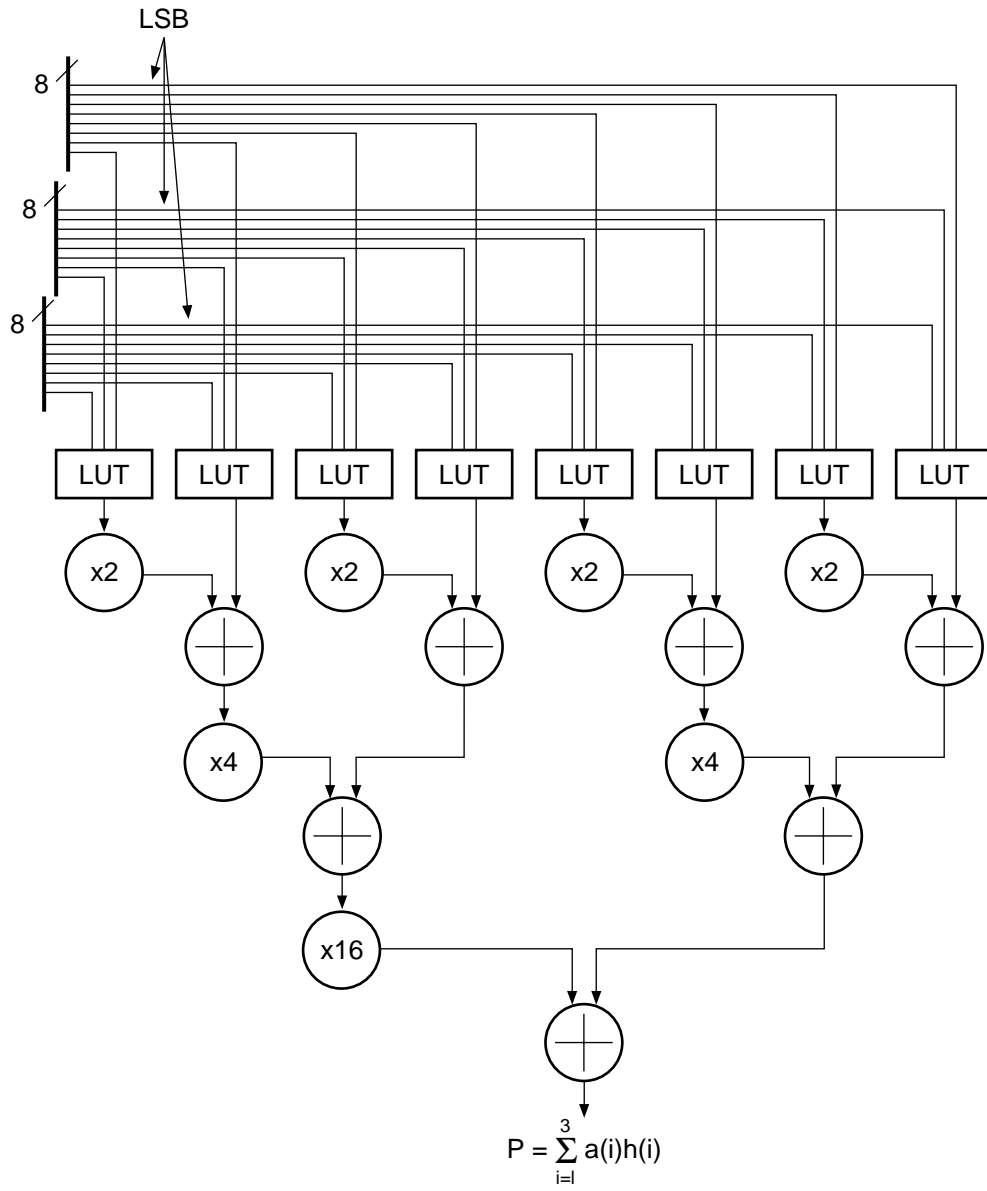
$$P1 = h(1) + h(2) + 0 = 11$$

The same procedure applies when forming P2. We look at the second LSBs of the $a(n)$. For the bit combination 011, we form:

$$P2 = 0 + h(2) + h(3) = 10$$

Finally, we arithmetic left-shift P2 by one bit and add to P1 to get $P=111$. Taking advantage of this combinatorial property of the input bits, we can pre-compute all possible combinations (see Table 1) and store them into a LUT. When performing vector multiplication, we then simply call the LUT for the specific input bit combination, and accumulate the LUT output with the appropriate arithmetic bit-shift(s). For reference, a three 8-bit input vector multiplier is shown in Figure 3.

Figure 3. Three 8-bit Input Vector Multiplier



Implementation

Our reference design assumes operation on 256-level grayscale images whose intensity values are unsigned 8-bit integers. Same assumption holds for the convolver mask taps for simplicity. The precision and the choice of number system are application specific, and are customizable with additional logic circuitry. An overall block diagram of our 3 x 3 symmetric convolver design is shown in Figure 4.

Input registers and Pre-addition

Three 8-bit parallel datastreams of the digitized image/video frame are clocked into the 3 x 3 registers every two global clock cycles (CLK_2). This clocking is essential for the TDM vector multiplier. Fully pipelined adders are used to pre-add inputs, as described above. The adders can either run on the global clock or the CLK_2 signal. All logic blocks are generated by Atmel IDS Component Generator.

Figure 4. Block Diagram for 3x3 Symmetric Convolver using TDM Vector Multiplier

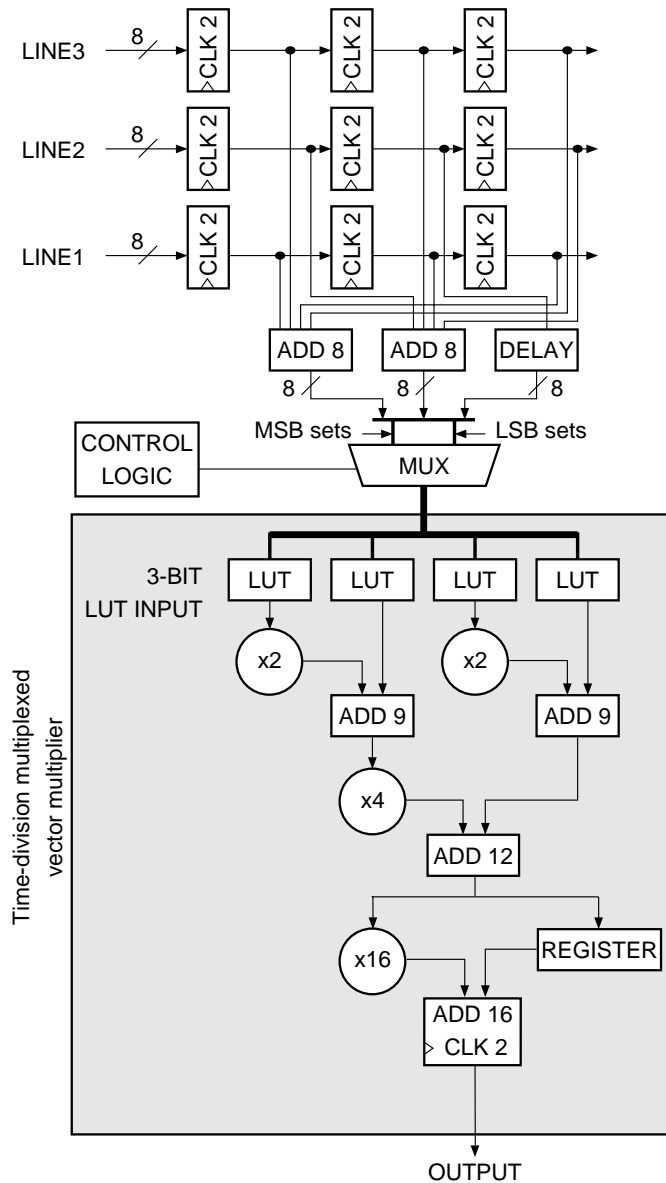
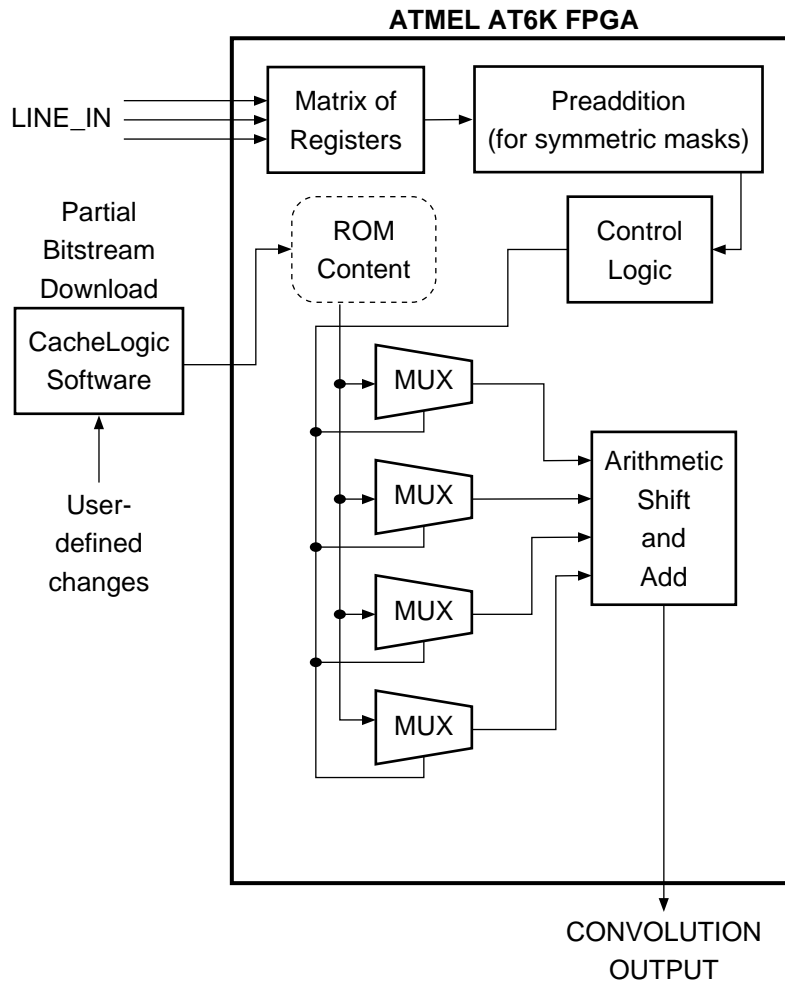


Figure 6. Architecture of TDM Convolver with Mux-constant ROM Emulation



SRAM implementation

As seen in Table 1, the LUT entries are simply the sums of different combinations of the constant mask coefficients. We can thereby configure our hardware to self-generate the LUT entries by using a counter to generate the combinations. Since we have three constant coefficients, we have a maximum of eight combinations of additions.

Figure 7 shows the LUT generator. A 3-bit ripple-carry counter is used to generate the combination signal. The AND gates are to “gate in” the coefficients to the adder stages. All adders are fully pipelined. The delay element is to compensate for the latency due to pipelining. Once the additions are performed, they will be stored in the four dual-port SRAMs that emulate the LUTs. Various control logic blocks will determine the activation of certain logic stages with precise timing. Note the address lines of the SRAMs. During the LUT configuration stage, the 3-bit ripple-carry counter will drive the address lines of the SRAMs. When the actual convolution is performed, the LSB and MSB combinations of the $a(n)$ will read the SRAMs for pre-computed partial products.

TDM vector multiplier

The TDM vector multiplier shown in Figure 4 operates on the same principle as the one shown in Figure 3. The only difference is the number of LUTs used. In our design, we split the look-up operation into two parts. We first operate on the four LSB sets of $a(n)$, performing the necessary shift and add. Then we repeat the same procedure for the remaining four MSB sets of $a(n)$. Notice the register immediately before the final adder stage. This register holds the result of the LSB operation so that the final addition is synchronized with the MSB sets of the same inputs $a(n)$. Again, this stage of the design is fully-pipelined, and runs on the global clock unless otherwise specified on Figure 4.

Figure 8. A Real-time Image Processing System

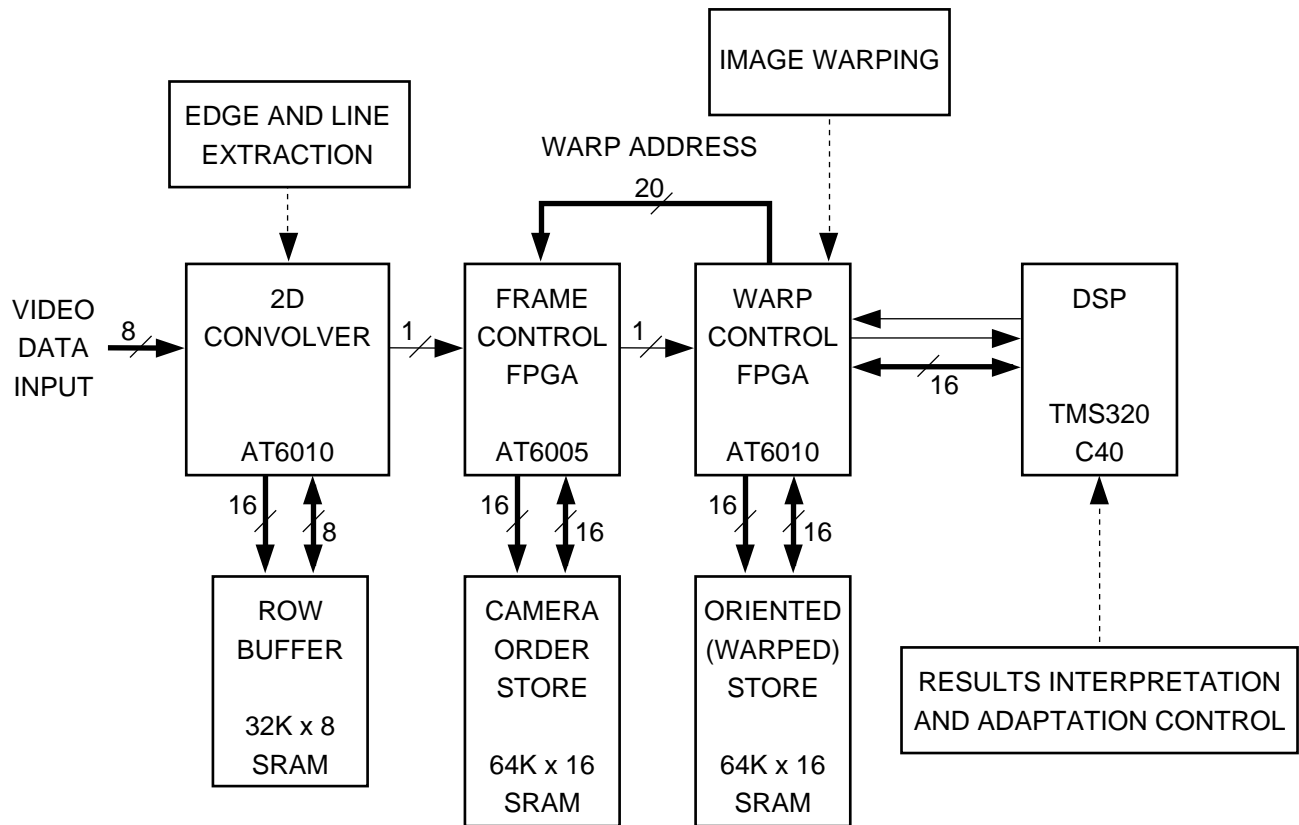


Table 2. Reference Design Statistics

LUT Implementation	ROM	RAM	Mux-Constant
Datapath	TDM Pipelined	TDM Pipelined	TDM Pipelined
Input Precision	8-bit	8-bit	8-bit
Partial Reconfiguration	CacheLogic	CacheLogic/LUT Generator	CacheLogic
Partial Reconfiguration Time	0.2 μ s per cell	0.2 μ s per cell	0.2 μ s per cell
Throughput Frequency (approx.)	15 MHz	15 MHz	25 MHz

References

- [1] Anil K. Jain, "Fundamentals of Digital Image Processing," Prentice-Hall, Inc., 1989.
- [2] Lee Ferguson, "Image Processing Using Reconfigurable FPGAs," DSP and Multimedia Technology, May/June 1996, Golden Gate Enterprises, Inc.



Atmel Headquarters

Corporate Headquarters

2325 Orchard Parkway
San Jose, CA 95131
TEL (408) 441-0311
FAX (408) 487-2600

Europe

Atmel U.K., Ltd.
Coliseum Business Centre
Riverside Way
Camberley, Surrey GU15 3YL
England
TEL (44) 1276-686-677
FAX (44) 1276-686-697

Asia

Atmel Asia, Ltd.
Room 1219
Chinachem Golden Plaza
77 Mody Road Tsimhatsui
East Kowloon
Hong Kong
TEL (852) 2721-9778
FAX (852) 2722-1369

Japan

Atmel Japan K.K.
9F, Tonetsu Shinkawa Bldg.
1-24-8 Shinkawa
Chuo-ku, Tokyo 104-0033
Japan
TEL (81) 3-3523-3551
FAX (81) 3-3523-7581

Atmel Operations

Atmel Colorado Springs

1150 E. Cheyenne Mtn. Blvd.
Colorado Springs, CO 80906
TEL (719) 576-3300
FAX (719) 540-1759

Atmel Rousset

Zone Industrielle
13106 Rousset Cedex
France
TEL (33) 4-4253-6000
FAX (33) 4-4253-6001

Fax-on-Demand

North America:
1-(800) 292-8635
International:
1-(408) 441-0732

e-mail

literature@atmel.com

Web Site

<http://www.atmel.com>

BBS

1-(408) 436-4309

© Atmel Corporation 1999.

Atmel Corporation makes no warranty for the use of its products, other than those expressly contained in the Company's standard warranty which is detailed in Atmel's Terms and Conditions located on the Company's web site. The Company assumes no responsibility for any errors which may appear in this document, reserves the right to change devices or specifications detailed herein at any time without notice, and does not make any commitment to update the information contained herein. No licenses to patents or other intellectual property of Atmel are granted by the Company in connection with the sale of Atmel products, expressly or by implication. Atmel's products are not authorized for use as critical components in life support devices or systems.

Marks bearing ® and/or ™ are registered trademarks and trademarks of Atmel Corporation.

Terms and product names in this document may be trademarks of others.



Printed on recycled paper.

0764B-08/99/xM